

## ROLE OF HIGH-END COMPUTING IN MEETING NASA'S SCIENCE AND ENGINEERING CHALLENGES

**Rupak Biswas, Eugene L. Tu, William R. Van Dalsem**  
NASA Ames Research Center, Moffett Field, CA 94035, USA  
{Rupak.Biswas, Eugene.L.Tu, William.R.VanDalsem}@nasa.gov

**Key words:** High-end computing, Columbia, computational fluid dynamics, integrated environment, parallel computing, aerospace applications, Ocean modeling, climate and weather modeling

### **Abstract**

High-end computing (HEC) plays a major role in the modeling and simulation needs of NASA missions. Two years ago, NASA was on the verge of dramatically increasing its HEC capability and capacity. With the 10,240-processor supercomputer, Columbia, now in production for 18 months, HEC has an even greater impact within the Agency and extending to partner institutions. Advanced science and engineering simulations in space exploration, shuttle operations, Earth sciences, and fundamental aeronautics research are occurring on Columbia, demonstrating its ability to accelerate NASA's exploration vision. This talk describes how the integrated production environment fostered at the NASA Advanced Supercomputing (NAS) facility at Ames Research Center is accelerating scientific discovery, achieving parametric analyses of multiple scenarios, and enhancing safety for NASA missions. We focus on Columbia's impact on two key engineering and science disciplines: Aerospace, and Climate. We also discuss future mission challenges and plans for NASA's next-generation HEC environment. (149 words)

### **1.0 Introduction**

High-end computing (HEC) has played a major role in meeting the modeling and simulation needs of NASA missions. Two years ago, having projected its near-term and future high-fidelity modeling and simulation requirements, NASA was on the verge of dramatically increasing its HEC capability and capacity. With NASA's 10,240-processor supercomputer, Columbia, now in production for 18 months, HEC is having an even greater impact within the Agency and extending to partner institutions. Significant cutting-edge science and engineering simulations in the areas of space exploration, shuttle operations, Earth sciences, and fundamental aeronautics research are occurring regularly on Columbia, demonstrating its ability to accelerate NASA's exploration vision. This talk describes how the integrated supercomputing production environment fostered at the NASA Advanced Supercomputing (NAS) facility located at Ames Research Center, is being used to develop aerospace vehicles, conduct parametric analysis for safe operation of the Shuttle, accelerate scientific discovery, and enhance safety during the life cycle of NASA missions. Columbia's impact is illustrated using two of the agency's key engineering and science disciplines: Aerospace and Climate/Weather.

In the Aerospace arena, computed results are presented in three areas: Debris transport analysis for the Space Shuttle's Return to Flight (RTF); Shuttle Main Engine flowliner analysis related to

mission safety; and risk assessment of ascent abort scenarios for proposed Crew Exploration Vehicle designs.

Two years ago, the role of Columbia's predecessor system in the Space Shuttle STS-107 accident investigation [1] was reported. Since that time, Columbia has been a key resource for the shuttle's Return to Flight effort and for damage assessment and repair recommendations for remaining shuttle flights. For future flights, including the launch scheduled for July 2006, CFD and Columbia play a large role, and simulations are integral to launch risk analysis and component redesign. Columbia is also having a significant impact on NASA's numerous space and exploration applications, such as the development of the Crew Exploration and Launch Vehicles (CEV/CLV), and risk assessment through the mission cycle—from ground operation, launch, and return to earth.

Among NASA's applications in climate and weather modeling are: next-generation global ocean models that resolve eddies and other narrow current systems, and atmospheric modeling and prediction of hurricane tracks for early warning.

The role of the Columbia supercomputer (currently ranked the fourth fastest system in the world, at 62 teraflops peak performance) in advancing the science and technologies related to the above topics are illustrated through various data analysis methods. Users of Columbia are also supported by the NAS facility's integrated HEC environment. In addition to system analysts, experts in code parallelization and performance optimization, high-fidelity modeling and simulation, high-speed networking, and data analysis and visualization exploit the power of Columbia to enhance NASA's computational capability.

As with other federal agencies and industry, NASA's future mission challenges require even more powerful systems. At present, the Agency is planning its next-generation HEC system environment, which is expected to be about four times more powerful than Columbia. Development of this future environment includes anticipated storage and archive requirements for a balanced system, application performance enhancement tools, and new wide area networking technologies.

## **2.0 Aerospace Applications**

### **2.1 Debris transport analysis for Space Shuttle launch scenarios**

After the STS-107 incident in February 2003, the Columbia Accident Investigation Board (CAIB) requested that CFD researchers conduct a debris transport analysis, which provided insight into actual mechanism of debris shedding from the bi-pod ramp region and provided input for test velocity and density for foam impact testing to determine damage to the orbiter's wing leading edge [2].

Subsequently researchers at NASA Ames have developed a CFD process for determining the aerodynamic characteristics of debris shedding during the Space Shuttle Launch Vehicle (SSLV)

ascent. Then, a complete debris scenario was conducted on Columbia, which focused on predicting the aerodynamic characteristics of potential debris sources, such as insulating foam and ice. This computational analysis was critical to NASA's Return to Flight effort, where scientists performed six-degree-of freedom (6-DOF) foam debris transport analyses and visualization to forecast shuttle damage, and for damage assessment and repairs during the successful Space Shuttle Discovery flight in July 2005.

For future flights, debris analysis has been directed to an assortment of shapes (typically thin and conical) that can potentially be shed from the external tank (ET) foam. The debris sources and their aerodynamic characteristics, are put into the debris transport code, which calculates trajectory information to assess the potential damage or risk by a specific debris source to a specific structural component—such as the ET foam impacting the orbiter wing. For this interactive process to be effective, the debris transport analysis must be done rapidly. A single trajectory calculation requires 30-60 CPU-hours on the Columbia system, which provides enough throughput to rapidly and efficiently run hundreds of trajectories in a day, using only a fraction of the computational resources. A system with the power of Columbia is required to run the typically hundreds of thousands of trajectories analyzed over the entire vehicle for each iteration.

The results in Figure 1 show that the average drag for the oscillating trajectory of an idealized frustum and the tumbling trajectory of a highly asymmetric debris piece are similar.

This is not the case when considering the crossrange behavior (Figure 2). The dynamically stable oscillating frustum generates almost no crossrange, as the lift force oscillates first in one direction then the other, with little net effect. In order to provide a representative distribution, researchers used several shapes to develop the crossrange constraints. These include real digitized shapes, idealized frustums, ideal frustums with the center of mass offset, and slightly asymmetric shapes such as elliptical frustums with the small diameter slightly offset from the larger.

The crossrange envelopes show a zero-lift trajectory emanating from the ET flange region. The intersection of this cone with the SSLV shows that the fuselage and wing of the orbiter have potential for debris impacts from this flange location, along with regions of the left solid rocket booster.

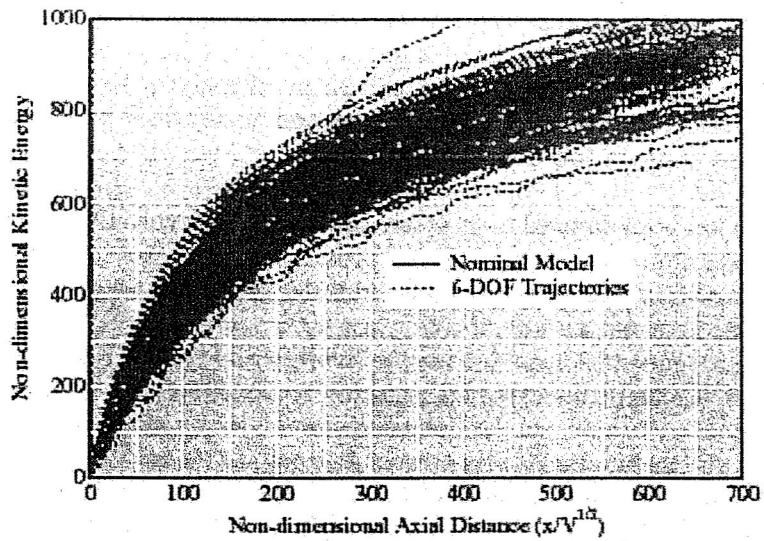


Figure 1: Comparison of kinetic energy variation for unconstrained 6-DOF trajectories and integrated trajectories using a nominal drag model based on the ensemble-average curves for a range of shapes. Trajectories released at Mach 2.5.

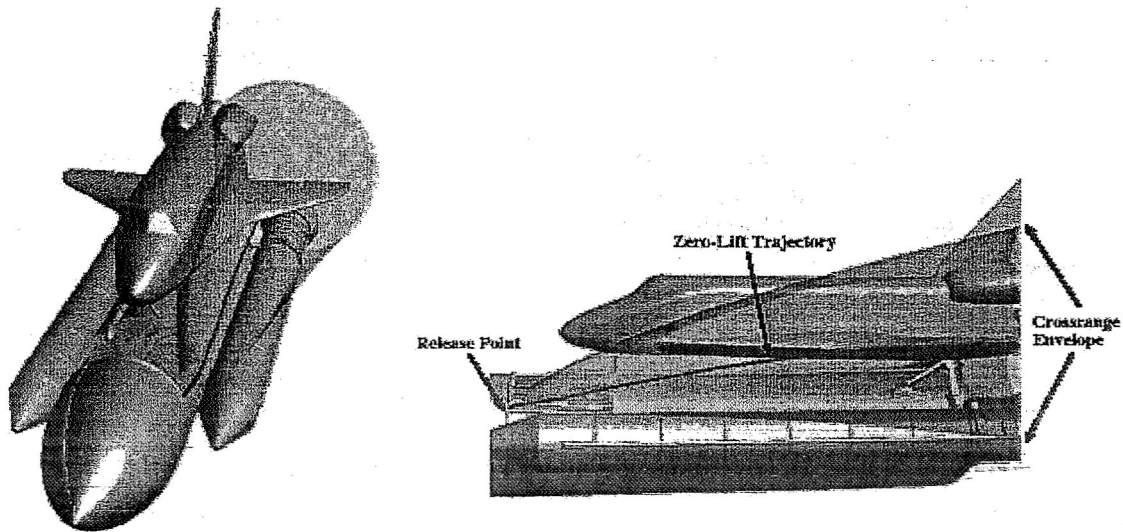


Figure 2: Crossrange envelope superimposed upon the computed ballistic zero-lift trajectory. A statistical distribution of the crossrange within the envelope can be used for a risk analysis.

Results show the farther downstream the debris travels before impact, the greater the impact kinetic energy, as the aerodynamic drag is constantly increasing the relative velocity between the debris and the orbiter.

The emphasis of this work is the development of an efficient process for modeling debris beyond ET insulating foam, including the ET LOX frost ramps, insulating cork on the solid-rocket boosters, frost and ice on the ET acreage regions, and ice which can form on the ET feedline brackets. The flexibility of the modeling and simulation capability and the computing resources provided by Columbia allows the dynamic behavior of these diverse debris sources to be analyzed in a systematic and timely manner.

## **2.2 Space Shuttle Main Engine Flowliner Analysis**

In May 2002, numerous cracks were found in the Space Shuttle Main Engine (SSME) #1 flowliner; specifically, at the gimbal joint in the liquid hydrogen (LH<sub>2</sub>) feedline flowliner. Since then, repairs have been made to existing cracks on all orbiters. Long-term scientific investigations continue, as the root cause of the original cracks was not conclusively established, and remaining shuttle flights are involved.

High-fidelity computations have been conducted on the Columbia supercomputer to analyze the SSME LH<sub>2</sub> feedline flowliner [3]. Numerous computational models were used to characterize the unsteady flow features in the turbopump, including the Low-Pressure-Fuel-Turbopump (LPFTP) inducer, the orbiter manifold, and an experimental hot fire test article representing the manifold. Findings show that unsteady flow stemming from the LPFTP inducer is one of the major contributors to high-frequency cyclic loading that results in fatigue damage to the flowliners.

The flow fields for the orbiter manifold and the hot fire test article were computed and analyzed on Columbia for similarities and differences using an incompressible Navier-Stokes flow solver, INS3D [4-6].

The first computational model included the LPFTP inducer; by studying the inducer model alone, scientists were able to compare unsteady pressure values against existing data. To resolve the complex geometry in relative motion, an overset grid approach [7] was employed, which contained 57 overlapping zones with 26.1 million grid points. The second computational grid system added the flowliner geometry. This grid system, which was very similar to the ground test article, consisted of 264 overset grids with 65.9 million grid points, and is shown in Figure 3. The flowliner component alone contained 212 grids and 41 million points.

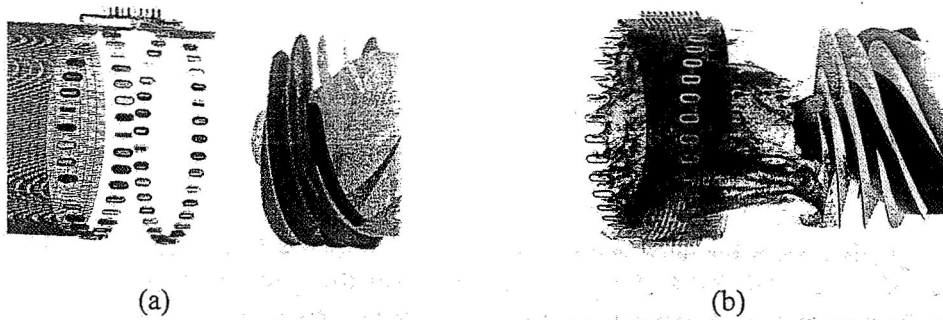


Figure 3. Computational model for LPFTP inducer and the liquid LH<sub>2</sub> flowliner: (a) Grid, (b) Computed results illustrating unsteady flow interaction of the flow in the bellows cavity and the back flow from the inducer.

To accelerate and automate the grid generation process, scripts were developed to automatically create grids for each type of component. The size of the simulation is large, requiring parallel processing to obtain solutions with reasonable turnaround. Two parallel programming paradigms were leveraged in the INS3D code: the Multi-Level Parallelism (MLP) [8] and the hybrid MPI+OpenMP models. Multiple-node computations showed that point-to-point implementation of the MPI+OpenMP code performs more efficiently than the master-worker version of the MPI code [9] (see Figure 4).

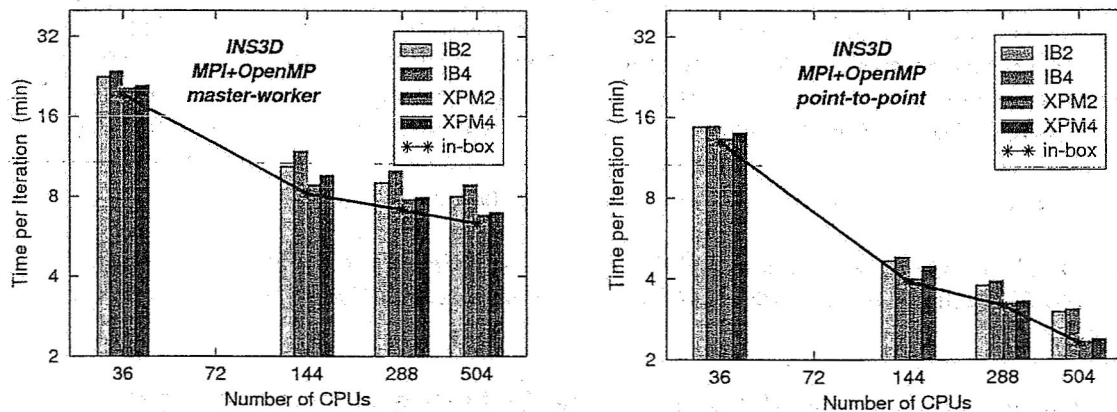


Figure 4. Multi-node INS3D performance.

Results of the CFD calculations confirmed the presence of backflow caused by the LPFTP inducer. The region of reverse flow extended far enough upstream to interfere with both flowliners in the gimbal joint. Computed results for the test article were verified by correlation with pressure measurements, and confirmed a strong unsteady interaction between the backflow caused by the LPFTP inducer and secondary flow in the bellows cavity through the flowliner slots. It was observed that a swirl on the duct side of the downstream flowliner is stronger than on the same side of the upstream flowliner—causing significantly stronger unsteady interactions

through the downstream slots than those observed in those upstream. This turbopump application currently exhibits some of the best scalability on the Columbia system.

### **2.3 Risk Assessment of Abort Scenarios for Proposed CEV Designs**

Researchers are running high-fidelity computational fluid dynamics (CFD) codes on Columbia for future space vehicle designs, including the Crew Exploration and Launch Vehicles (CEV/CLV), and are building realistic models to simulate flight risks for these new spacecraft. Risks and performance issues during both the ascent and entry/descent/landing phases are being carefully analyzed. The CEV will transport a maximum of six crew members to and from the International Space Station and up to four astronauts to and from the moon.

The CEV design includes a Launch Abort System (LAS) for crew escape, similar to that used in the Apollo capsule. Several computational modeling and simulation tools suited for analyzing abort scenarios have recently been developed and enhanced for use on Columbia. Staff at NASA's Ames and Glenn Research Centers have collaborated on this work under the Simulation Assisted Risk Assessment (SARA) project. The SARA team developed a Probabilistic Risk Assessment (PRA) approach and demonstrated how risk analysis can be applied to launch abort using the Apollo configuration [10]. A PRA identifies the best level of fidelity for modeling critical failure modes associated with launch abort. Columbia was then used to conduct higher-fidelity modeling on specific failure modes. Two failure modes examined included booster explosion and those caused by re-contact with the booster during separation. Each of these modes required the application of high-fidelity aerodynamic simulation.

Analysis of the booster failure mode (using Apollo data), showed a possible catastrophic failure of the booster, leading to detonation of the propellant, which creates blast wave overpressures that could fatally damage the LAS (Figure 5). As the risk model was being developed, it became clear that the type of booster and the nature of the failure it was likely to encounter determined the environments under which the crew escape system must operate to ensure a successful abort. The process for characterizing this interaction must be carefully modeled and simulated.

One particular weakness found in an engineering-level model was the effect of headwind as the CEV ascends. To account for these effects in the risk analysis, high-fidelity blast wave models were built and simulated on Columbia using the Overflow Navier-Stokes code [11]. Results indicated that headwinds significantly affect the nature and magnitude of the shock wave as it impacts an escaping CEV. This means that the warning time required to initiate the abort sequence is also affected. Additional work in high-fidelity simulations is being done to help engineers generate requirements for the LAS.

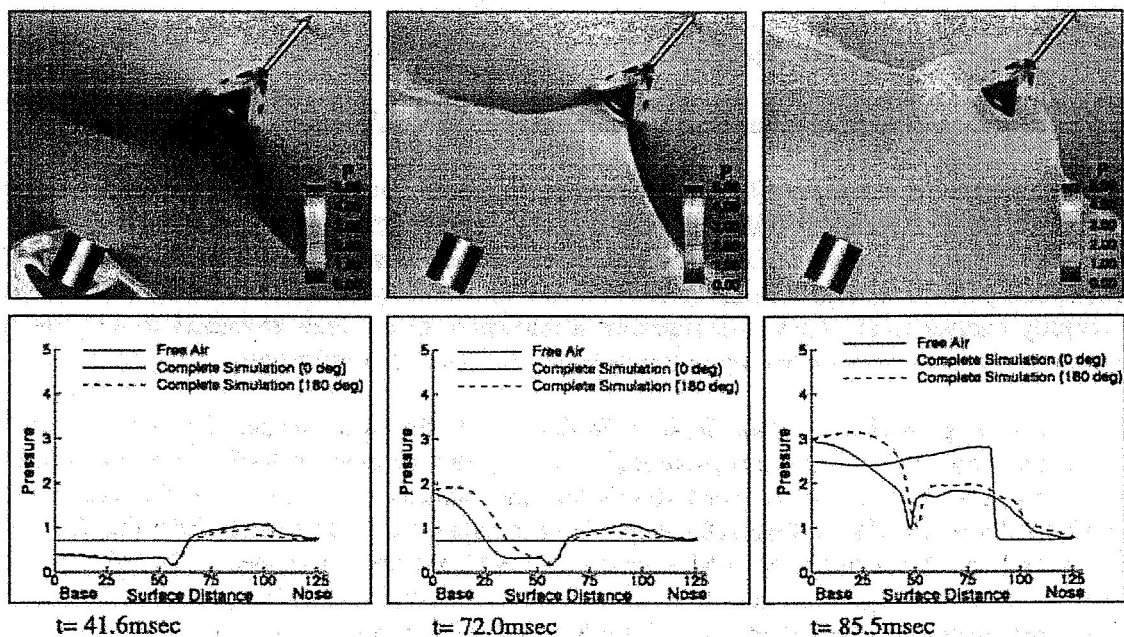


Figure 5. Flowfield and surface pressures for blast wave propagating through wake of maneuvering Launch Abort System.

Another failure mode dependent on high-fidelity simulation involves the ability of the LAS to achieve “clean” separation of the CEV from the booster stack in the event of impending catastrophic failure; simply put, the CEV must not scrape or re-contact the booster stack. This failure mode was especially demanding because it involved complex proximity aerodynamics—modeling transonic flow and the complex flow at the small gap (or cavity) between the CEV and booster stack at separation. Both Navier-Stokes simulations, using Overflow, and Euler simulations, using FlowCart [12] were applied, and their results validated against transonic wind tunnel and abort flight test data from the Apollo era [13].

All these cases are computationally expensive to simulate. A single steady-state simulation required approximately 3,500 processor-hours on Columbia. The complexity of the geometry and the flow-field required about 30 million grid points, which enabled good scalable performance up to at least 250 CPUs. In all, approximately 20 cases were computed using Overflow at various ascent trajectories and separation thrust levels. Each case required approximately 20,000 processor-hours on Columbia. All failure modes benefited greatly from the HEC resources at the NAS facility. These tools and processes, utilizing the Columbia resources, will likely be applied to analyze the actual LAS design, and to further understand the CEV failure modes and their impact on the vehicle’s survivability.

### 3.0 Climate and Weather Applications

#### 3.1 Global Modeling of Ocean Currents

To increase understanding and predictive capability for the ocean's role in future climate change scenarios, NASA has begun a project called Estimating the Circulation and Climate of the Ocean, Phase II (ECCO2): High-Resolution Global-Ocean and Sea-Ice Data Synthesis [14]. The goal is to produce increasingly accurate combinations of all available global-scale ocean and sea-ice data at resolutions that begin to resolve ocean eddies and other narrow current systems, which transport heat, carbon, and other properties in the ocean. Among other scientific applications, these data syntheses are used to quantify the role of the oceans in the Earth's carbon cycle, understand recent changes in the polar oceans, and monitor time-evolving term balances within and between different components of the Earth system. This work aims to harness NASA's computational resources, advances in CFD and software engineering, and the ability to solve massive control problems.

The most challenging numerical experiment undertaken to date is a near-global simulation with  $1/16$  horizontal grid spacing, that is, approximately 6 km at the Equator and 1 km at high latitudes. The number of surface grid cells is just over 25 million and the configuration has 50 vertical levels, bringing the total number of cells in all three dimensions to just over 1.25 billion. Each of the 3-D fields that describe the simulation domain and its time evolving state requires 10 gigabytes (GB) of storage. This configuration has been integrated on the 2,048-CPU cluster on Columbia [14] (see Section 4.1). This workload is spread evenly over 1,920 of Columbia's processors, so that each individual CPU is responsible for simulating about 586,000 grid cells (equivalent to a surface region roughly  $210 \times 210$  km). Decomposing the workload over this many processors yields a setup that, with extensive diagnostics and analysis options included, uses about 870 megabytes (MB) of main memory per processor. With a timestep of two minutes, this performance allows a year of simulation to be completed in less than ten days.

To investigate solution convergence as horizontal resolution is increased ECCO2, researchers have begun a series of numerical simulations at  $1/4^\circ$ ,  $1/8^\circ$ , and  $1/16^\circ$  resolutions. Figures 6 shows significant changes in solution with resolution. Each plot shows the change in simulated sea-surface height over the same month. The plots capture changes due to eddy activity over a single month. Changes with resolution occur in the regions of the oceans where eddies are prevalent (such as the Gulf Stream, the Kuroshio, the Agulhas, the Drake Passage and the Antarctic Circumpolar Current). For example, in the Gulf Stream region, a relatively small area of vigorous sea-surface height changes. Key behaviors, such as how tightly waters "stick" to the coast, or how far energetic eddies penetrate the ocean interior, change significantly between resolutions.

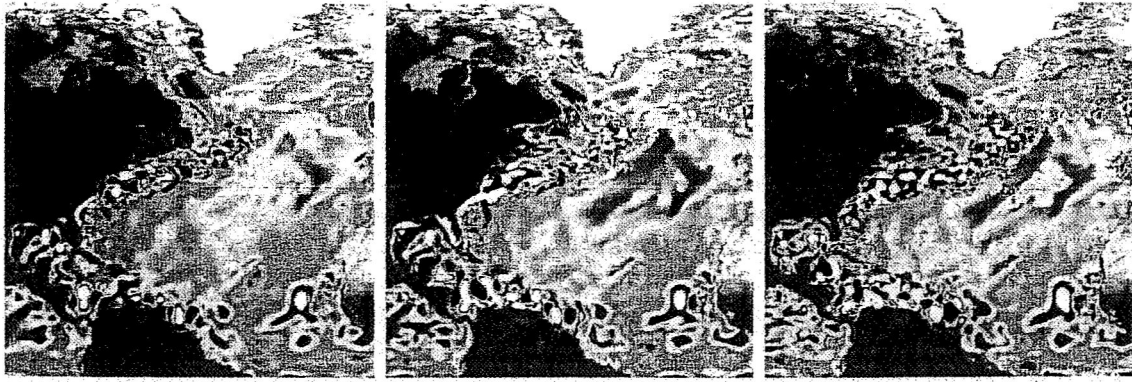


Figure 6: Gulf Stream region sea-surface height difference plots at different resolutions. Left panel  $1/4^\circ$ , middle panel  $1/8^\circ$ , right panel  $1/16^\circ$ . Color scale -0.125m to 0.125m.

Performance on the Altix shows that it is well suited for addressing these questions.

NASA's ECCO2 code achieves about 722 Mflop/s per CPU when running on 1,920 processors—14 percent of the per-CPU Top500 number achieved on the system [16]. The code consists of predominantly BLAS1 class operations and cannot exploit the level of cache reuse that the Top500 Linpack benchmark achieves. The scaling across multiple Altix systems is encouraging and suggests that configurations spanning eight or more Altix nodes—and that would therefore support  $1/20$  and higher resolutions—are within reach.

### 3.2 Atmospheric Modeling and Prediction of Hurricanes

The NASA Finite Volume General Circulation Model (fvGCM) is a unified numerical weather prediction (NWP) and climate model that could run on daily, monthly, decadal, and century time-scales. The model was originally designed for climate studies at a coarse resolution of about  $2 \times 2.5^\circ$ , and has been running at much finer resolution on Columbia to answer the following question for NASA's mission in hurricane research: how can weather/hurricane forecasts can be improved and made more reliable over longer periods of time using computer modeling? [17].

Hurricane forecasts pose challenges for General Circulation Model (GCMs), the most important being horizontal grid spacing. With the unique computing resources of the Columbia supercomputer, the model horizontal resolution was rapidly increased to  $1/4^\circ$  in 2004 e.g., [18], and  $1/8^\circ$  in early 2005. Recently, researchers have tested a  $1/12^\circ$  resolution version, which is the first global weather model with a single-digit resolution (9km at the equator). A five-day forecast of total precipitable water with the  $1/12$  degree fvGCM (Figure 7) shows fine scale weather events in the tropical area, which brings researchers to overcoming the fundamental barrier between global and mesoscale models [20].

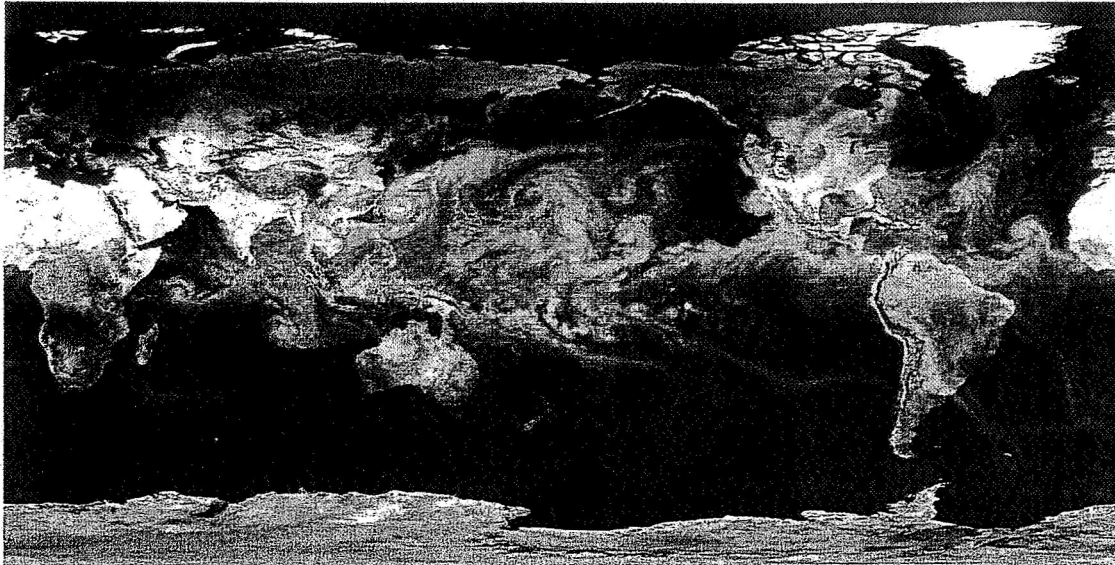


Figure 7. Five-day forecasts of total precipitable water initialized in September 2004 with the  $1/12^\circ$  fvGCM, giving a grid spacing of 9km in the equator.

During the 2004 hurricane season, the  $1/4^\circ$  model, which doubled the resolution adopted by most global models in operational NWP centers at that time, was running in real-time, and provided remarkable landfall predictions up to five days in advance for major hurricane such as Charley, Frances, Ivan, and Jeanne e.g., [19]. Moreover, the model proved capable of resolving problems such as erratic track, abrupt recurvature, and intense extratropical transition. In the 2005 hurricane season, new research focused on the validations of the  $1/8^\circ$  fvGCM performance on hurricane forecasts, while the real-time  $1/4^\circ$  forecasts provided a baseline for comparisons. Being a global mesoscale-resolving model, the  $1/8^\circ$  model was the first global model to simulate mesoscale vortices (such as the Catalina Eddy and the Hawaiian Lee Vortex), which were generated by the interaction of the large-scale flows with better resolved surface forcing.

The 2005 Atlantic hurricane season was the most active in recorded history. There were 28 tropical storms and 15 hurricanes, four of which were Category 5 hurricanes. Accurate forecasts of these storms brought a great challenge to global and mesoscale modelers. It is well known that GCMs' insufficient resolutions undermine intensity predictions. Using the power of Columbia, NASA researchers demonstrated that this limitation could be overcome by performing six 5-day forecasts of hurricane Katrina [21] with the  $1/8^\circ$  fvGCM, showing promising intensity forecasts with small errors in center pressure of only  $\pm 12$  hPa. Notable improvement in Katrina's intensity forecasts occurred when the grid spacing decreased from  $1/4^\circ$  to the  $1/8^\circ$ , at which the near-eye wind distribution and the radius of maximum wind could be resolved more realistically. While the mesoscale-resolving fvGCM has produced very promising results for the past two years, a great potential for further advancement is still ahead.

#### 4.0 Role of HEC

High end computing has played a key role in NASA missions for over 25 years. The Columbia supercomputer has increased the space agency's computing capability ten-fold and revitalized its high-end computing (HEC) effort. Constructed in just four months, Columbia has enabled U.S. scientists and engineers to perform important, breakthrough simulations since it became fully operational in October 2004. Performing at a peak of 62 teraflops (Tflop/s), Columbia has demonstrated its capability to support and accelerate the agency's key missions and vision for exploration.

#### 4.1 Columbia Description

Columbia is a 10,240-processor Linux-based SGI Altix cluster composed of twenty 512-processor nodes, twelve of which are SGI Altix 3700s, and the remaining eight are double-density SGI Altix 3700 Bx2 nodes. Each node is a shared memory, single-system-image (SSI) system, running the Linux operating system. Four of the Bx2 nodes are tightly linked to form a 2048-processor shared memory environment.

Each processor in the 2048-CPU subsystem is an Intel Itanium 2, running at 1.6 gigahertz (GHz), with 9 MB of level 3 cache (the "Madison 9M" processor), and a peak performance of 6.4 gigaflops (Gflop/s), with a total of 4 terabytes (TB) of shared memory, or 2 GB per processor. One other Bx2 node is equipped with these same processors. The remaining 15 nodes have Itanium 2 processors running at 1.5 GHz, with 6 MB of level 3 cache, and a peak performance of 6.0 Gflop/s. All these nodes also have 2 GB of shared memory per processor.

Within each node of Columbia, the processors are interconnected via SGI's proprietary NUMalink fabric. The 3700's utilize NUMalink3 with a peak bandwidth of 3.2 GB/s. The Bx2s have NUMalink4 where the bandwidth is doubled to 6.4 GB/s. The 20 nodes are connected to each other by Voltaire InfiniBand fabric, as well as via 10- and 1-gigabit Ethernet connections. The four Bx2 nodes in the 2048-CPU subsystem use NUMalink4 among themselves as well as the other fabrics. Columbia is connected to 440 TB of online RAID storage through a Fibre Channel switch. This capacity is being upgraded to almost 1 petabyte (PB). The archive (tape) storage capacity is 10 PB.

On each 512-processor node, the primary features are as follows:

- Low latency to memory (less than 1 microsecond), which significantly reduces the communication overhead;
- High memory bisection bandwidth, Columbia being the first system (in November 2004) to exceed 1 TB/s on the STREAM benchmark [2];
- Global shared memory and cache-coherency, which enables simpler and more efficient programming paradigms than message passing;
- Large shared memory (1 TB), which allows bigger problems to remain resident on the system.
- These features make Columbia particularly well suited for large-scale compute- and data-intensive applications. Typical problems are physics-based simulations involving a discretized grid of the physical domain that is partitioned across multiple processors. In addition, applications requiring dynamic load balancing and/or adaptive gridding are much easier to control on Columbia, leveraging shared memory programming models such as OpenMP and MLP [8].

The development and operating environment on Columbia features a 64-processor SGI Altix front-end, a Linux-based operating system, Altair PBS Professional job scheduler, Intel Fortran/C/C++ compiler, and SGI ProPack software.

## 5.0 Vision for the Future and Concluding Remarks

Simulation methodologies have advanced along with computational technologies in general. Advanced tools have been developed to the point that many daily engineering and science problems can now be routinely computed; however, this is still done mostly using geometrically and or physically simplified or truncated models. Some of the physical models, such as those for turbulence and transition, and for high-temperature real gas, have not been advanced much more than what was available in the 1970s or '80s.

To realize the full benefit of HEC, more inclusive modeling of geometry and physics is needed. Attempts to solve these problems have been made with some qualitative success. However, predictive capability is still very limited and prediction with accurate physics is yet to be accomplished; this will require inclusion of not only fluid dynamic quantities but other quantities like thermal loading, structural properties, and control. These computations will require not only large computing resources but also increased data storage and sophisticated management technologies.

Many of Columbia's scientific and engineering users have stated that the system has allowed them to successfully complete investigations they never allowed themselves to dream of previously. Now, these users are envisioning what they can accomplish when even more powerful computing systems are available. NASA and the HEC community are working on developing petaflop computers that can execute  $10^{15}$  operations per second. The National Science Foundation (NSF) has a plan to provide computational resources to the community that can sustain petaflop performance by 2010. NASA is currently planning its next-generation supercomputer to meet the ever-increasing demand for computational resources required for a wide range of scientific discoveries and engineering applications.

For example, with NASA's next-generation system, scientists envision a launch simulation model designed to treat the whole launch environment until the vehicle has cleared the launch tower. The model would integrate 6-DOF multiple-body motion, debris impact, propulsion system vibration and exhaust, acoustics due to exhaust, fuel accumulation in the exhaust plume, exhaust chemistry including fuel burning, thermal stress on the vehicle structure and finally weather at the launch site. This very complex model would integrate data from propulsion simulation, meso-scale weather prediction and experiment. Utilizing state-of-the-art flow simulation tools and a next generation computing system, researchers can attempt to compute 100-million plus point grid aerodynamic simulation within a realistic turnaround time—within days rather than weeks.

In aerospace design, the most productive aspect of HEC applications has been to predict relative change among design variations. To push the limit of operation and to try bold new ideas, more predictive capabilities will be needed for complicated physical phenomena. Without accurate prediction, the capability impacts of HEC can be limited to the current level, even if more advanced facilities become available. To make these advances, high-fidelity computation using high-end computing facilities will still be a "must" despite all the excitement about the PC clusters and grid computing.

A final note is related to human resources. Although modeling and simulation technology has advanced remarkably, many challenging cases require experts in computational physics. Computer science can automate a good portion of the simulation processes, thus saving a large amount of the human effort required. However, blind application of tools without understanding capabilities and limitations of the methods involved could lead to catastrophic engineering results. As in many other engineering and science disciplines, modeling and simulation researchers and practitioners need to understand physics and the engineering systems being simulated. Experts who are willing to think through the flow physics in addition to software engineering, must still be developed for future generations.

## 7.0 References

- [1] F. Ronald Bailey, "High-End Computing Challenges in Aerospace Design and Engineering," International Conference on Computational Fluid Dynamics 3, Toronto, July 12–16, 2004.
- [2] Scott M. Murman Michael J. Aftosmis and Stuart E. Rogers Characterization of Space Shuttle Ascent Debris Aerodynamics Using CFD Methods 43rd AIAA Aerospace Reno, NV, January 10–13, 2005.
- [3] Kiris, C., Kwak, D., Chan, W., and Housman, J 'High-Fidelity simulations for Unsteady Flow Through Turbopumps and Flowliners,' *AIAA Paper 2006-0089*, AIAA 44th Aerospace Sciences Meeting, Reno, Nevada, Jan. 9-12, 2006.
- [4] Kwak, D., Chang, J. L. C., Shanks, S. P., and Chakravarthy, S., 'An Incompressible Navier-Stokes Flow Solver in Three-Dimensional Curvilinear Coordinate Systems Using Primitive Variables,' *AIAA J.*, vol. 24, No. 3, 390-396, Mar. 1986.
- [5] Kiris, C., and Kwak, D., 'Numerical Solution of Incompressible Navier-Stokes Equations Using a Fractional-Step Approach,' *Computers & Fluids*, vol. 30, 2001, pp829-851.
- [6] C. Kiris, D. Kwak, and W. Chan, "Parallel unsteady turbopump simulations for liquid rocket engines," *Supercomputing2000*, 2000.
- [7] P.G. Buning et al., "Overflow user's manual, version 1.8g", *Technical Report*, NASA Langley Research Center, 1999.
- [8] J.R. Taft, "Achieving 60 Gflop/s on the production CFD code Overflow-MLP", *Parallel Computing*, 27, 2001, pp. 521-536.
- [9] R. Biswas et al., "Performance characterization of the Columbia supercomputer using scientific applications", *Supercomputing2005*, 2005.
- [10] S. Lawrence et al, "Simulation assisted risk assessment," *44th AIAA Aerospace Sciences Meeting*, 2006, Paper 2006-0090.
- [11] S. Lawrence, D.L. Mathias, K. Gee, and M.E. Olsen, "Simulation assisted risk assessment: Blast overpressure modeling," *8th International Conference on Probabilistic Safety Assessment and Management*, 2006, Paper 0197.
- [12] M.J. Aftosmis, M.J. Berger, and G. Adomavicius, "A parallel multilevel method for adaptively refined Cartesian grids with embedded boundaries," *38th AIAA Aerospace Sciences Meeting*, 2000, Paper 2000-0808.
- [13] B.L. Barrier and O.C. Pendergraft Jr., "Transonic aerodynamic characteristics of a powered wind-tunnel model of the Apollo Launch Escape Vehicle during separation," NASA TM-X 1336, 1967.
- [14] Dimitris Menemenlis, et al., "NASA supercomputer improves prospects for ocean climate research," *Eos Trans. AGU*, vol. 86, pp 89, 95-96, 2005.
- [15] Chris Hill, Dimitris Menemenlis, Bob Ciotti, and Chris Henze, "Investigating solution convergence in a global ocean model using a 2048-processor cluster of distributed shared

- memory machine,” submitted to J. Scientific Programming (available at <http://ecco.jpl.nasa.gov/~dimitri/articles/HillJSciComp06.pdf>).
- [16] H. Meuer and J. Dongarra. The Top 500 benchmark web site. <http://www.top500.org>, 2005.
  - [17] NASA’s Role in Hurricane Research.  
[http://www.nasa.gov/pdf/147432main\\_hurr\\_fact\\_sheet.pdf](http://www.nasa.gov/pdf/147432main_hurr_fact_sheet.pdf)
  - [18] Davis, C. and L. F. Bosart, 2002: Numerical Simulations of the Genesis of Hurricane Diana (1984). Part II: Sensitivity of Track and Intensity Prediction. *Mon. Wea. Rev.*, **130**, 1100-1130.
  - [19] Atlas, R., O. Reale, B.-W. Shen, S.-J. Lin, J.-D. Chern, W. Putman, T. Lee, K.-S. Yeh, M. Bosilovich, and J. Radakovich, 2005: Hurricane forecasting with the high-resolution NASA finite volume general circulation model. *Geophys. Res. Lett.*, **32**, L03801, doi:10.1029/2004GL021513.
  - [20] Shen, B.-W., R. Atlas, J.-D. Chern, O. Reale, S.-J. Lin, T. Lee, and J. Chang, 2006a: The 0.125 degree finite-volume General Circulation Model on the NASA Columbia supercomputer: preliminary simulations of mesoscale vortices. *Geophys. Res. Lett.*, 2006GL026143, in press.
  - [21] Shen, B.-W., R. Atlas, O. Reale, S.-J. Lin, J.-D. Chern, J. Chang, C. Henze, 2006b: Simulations of Hurricane Katrina (2005) with the 0.125 degree finite-volume General Circulation Model on the NASA Columbia Supercomputer. *Geophys. Res. Lett.*, **33**, L05801, doi:10.1029/2005GL024594.